Operations Measurements for Engineering Support of High-Speed Networks with Self-Similar Traffic

J.L. Jerkins[a] A.L. Neidhardt[b] J.L. Wang[c] A. Erramilli[d]

[a]Bellcore, 331 Newman Springs Road, Red Bank, NJ 07701, USA, jlj@research.bellcore.com

[b]Bellcore, 331 Newman Springs Road, Red Bank, NJ 07701, USA, arnie@research.bellcore.com

[c]Bellcore, 331 Newman Springs Road, Red Bank, NJ 07701, USA, jwang@research.bellcore.com

[d]Qnetworx, 1119 Campus Drive West, Morganville, NJ 07751, USA, ashok@qnetworx.com

Traffic engineering for high-speed packet networks is necessary to ensure that provisioned capacity is sufficient to meet anticipated demand for a broad range of offered services, but not so excessive as to render services uneconomical. Currently most packet-based networks, including Internet, Frame Relay, and ATM networks, are engineered and managed either using traditional circuit-switched methods or on an ad-hoc basis. Given the anticipated diversity of packet applications (e.g., voice, data, and video) with fundamental differences in traffic patterns, usage, and QOS objectives, lack of more suitable methods can result in capacity and performance problems. Compounding the problem is the limited engineering support sometimes provided by vendors of packet network switches/routers, both in terms of operations measurements and traffic-related guidelines.

A fundamental issue of high-speed network engineering is the limited set of available *traffic measurements* which form the basis of guidelines that can be applied in practice. Using ATM as an example, we review some operational traffic measurements, including those in Bellcore generic requirements (GR-1248-CORE) for ATM operations, as well as propose measurement alternatives such as finer time-scale cell counts, peak rate measurements, link and processor utilization histograms, buffer usage measurements, and variance-time measurements, and describe how these measurements can be used in engineering. One or more of these new alternatives can potentially be implemented by the switch/router or management system suppliers to fulfill the need to better engineer high-speed packet networks.

1. INTRODUCTION

High-speed network services, for example, Internet, Frame Relay and ATM, have recently enjoyed tremendous growth; this in itself will not assure the profitability of these services. It is necessary to support this rapid growth with sound network traffic management practices so as to utilize existing capacity adequately without jeopardizing service integrity or profitability. Recent measurement results based on high-resolution traffic traces collected from working packet networks and services have shown that high-speed packet traffic exhibits fractal properties (see for example, [5–7,13,15,16,19,20,25]), which are fundamentally different from features found in circuit switched voice traffic and captured by commonly used packet traffic models such as batch-Poisson and Markov Modulated Poisson Process (MMPP). These fractal properties, *self-similarity* and *long-range dependence*, are associated with the well-known *burstiness* of packet traffic. There is no typical or characteristic value for traffic processes such as burst lengths, connection holding times, durations of activity and inactivity. Instead, these can span many time scales, referred to as infinite variance or the *Noah Effect*. Correlations in bursty traffic can also span many time scales, referred to as long-range dependence or the *Joseph Effect*. Both these features observed in real traffic can have significant impacts on traffic performance and engineering [9,12].

The first three columns of Figure 1 give time series plots of the aggregate traffic collected from three different network technologies: 1.5 Mbps Frame Relay (left), 10 Mbps Ethernet (middle), and 155 Mbps ATM (right). Given are the number of frames or cells transmitted over a link or trunk for four different time scales: 10 seconds (top), 1 second (second), 100 milliseconds (third), and 10 milliseconds (bottom).

Subintervals viewed on a smaller time scale are indicated by a darker shade in each plot. The traffic exhibits variations over four decades of time scales (i.e., the traffic does not appear to "smooth" to a fairly constant level), an indication of its fractal nature. These columns can be compared to the fourth column of Figure 1, which shows a generated sample of Fractional Brownian Motion [24], a model often used to describe certain types of data traffic (see details in the following section).

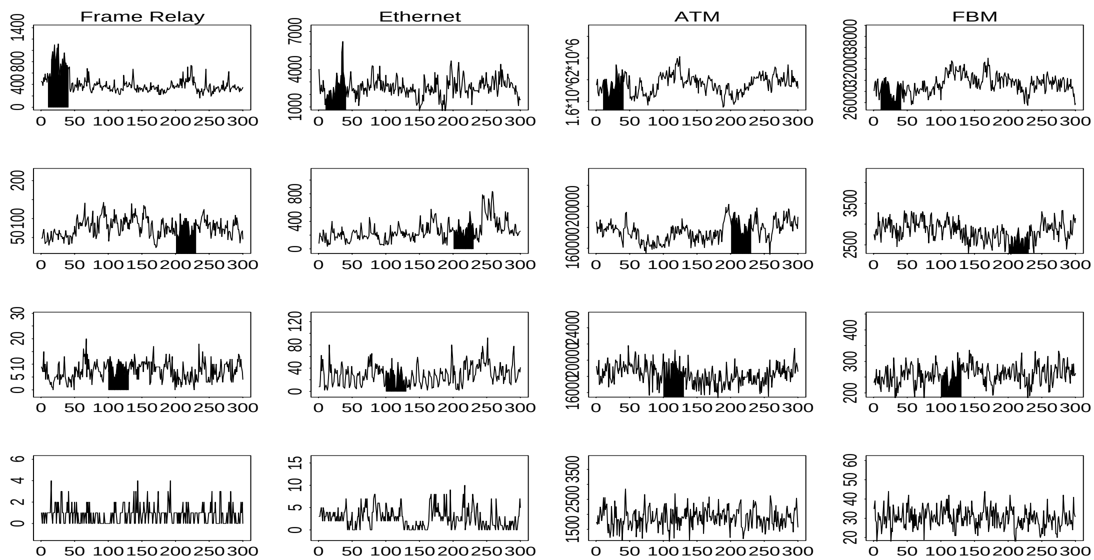

Figure 1. Variations of frame or cell counts over many time scales

Despite invaluable insights into the nature of packet traffic and resource usage gained through such aforementioned high-resolution studies, such collections and analyses on an on-going basis are not economically feasible. A more appropriate solution is for the network element (e.g., ATM switch or IP router) to collect measurements for on-line or off-line traffic parameter estimation and performance evaluation. [1]

Much like the UPCO (Usage, Peg Count, and Overflow) measurements collected by voice circuit switches, operational traffic measurements for high-speed networks are required for appropriately managing the underlying networks on a regular basis; good traffic engineering methods rely on parameterized traffic models where the parameters are estimated on the basis of network measurements. Detailed measurements would improve the accuracy and robustness of parameter estimation, but collection, handling, storage and analysis of many measurements can become a network bottleneck on its own.

In this paper, we review a number of operational traffic measurements, including those specified in Bellcore generic requirements for ATM (GR-1248-CORE [14]), as well as propose measurement alternatives which could enhance our ability to manage high-speed networks. *Note that the rest of this paper uses mostly ATM terminologies and measurement requirements as examples; the measurements discussed should apply to other packet technologies as well.* In the next section, we discuss the importance of traffic measurements, describe various ways of collecting them, and emphasize the need for a *parsimonious* traffic model supported by operations measurements. Section 3 briefly reviews traffic measurement requirements as specified in GR-1248-CORE. In Section 4 we describe various measurement alternatives and how they can be used in ATM traffic engineering, including finer time-scale cell counts, variance-time measurements, peak rate measurements, link and processor utilization histograms, and buffer usage measurements. A summary is provided in Section 5. The Appendix gives derivations of adjustments to the

[1] Cisco's NetFlow [21] is an example of such a set of measurements in the data network switch/router environment. See [3] for a development effort of `cflowd` which provides analysis and reporting capabilities based on NetFlow data.

variance-time analysis [19] needed because of the dependence in the traffic measurements; this discussion is required to facilitate our analyses of some of the measurement alternatives proposed, in particular, the special studies cell counts (Section 4.1) and variance-time measurements (Section 4.2).

2. TRAFFIC MEASUREMENTS

Traffic measurements form the basis of any engineering methodology. In principle, one can develop engineering guidelines on the basis of accurate and tractable representations of traffic, but unless supporting measurements are available, the guidelines are unusable in practice. In general, three types of traffic measurements are feasible: those *regularly scheduled* and collected by a switch, those collected as part of *special studies by a switch* or adjunct measurement device, and *high-resolution measurements* that require highly specialized traffic collection and storage capabilities. For the first two types of measurements, the reporting can be done by the switch sending the measurements to downstream data collection systems autonomously (after advance setup) or based on polling (e.g., using SNMP queries) through external data collection devices. These two types are specified in [14] as part of the broadband switching system data collection requirements. Most existing switches have only the first type of collection capability where the switch reports certain cell counts over a coarse time scale period, e.g., 15 minutes. The third type of data collection, high-resolution measurement of carried traffic, is necessary to explore and analyze ATM traffic patterns and resource usage to support other traffic management functions, see for example, [16,17].

In voice telephony, peg count is used to estimate average call rate, the *only* parameter required in the *Poisson* voice traffic model assuming average call holding time is known. Requirements for an appropriate packet traffic model include adequately capturing characteristics with performance impact, and also require a model with as few parameters as possible, essential for minimal overhead in processing. (The main factor limiting the use of many sophisticated packet traffic theories in the actual engineering process today is not tractability or motivation, but difficulties in setting model parameters [8]; models with more than a few parameters simply cannot be supported in practice.) To achieve parsimony as well as statistically describe packet traffic, the preferred approach is to model the variation of burstiness over many time scales parametrically instead of on each time scale explicitly (which would result in many parameters).

One such parsimonious model is the Fractional Brownian Motion (FBM) model, incorporating the self-similarity observed in actual traffic with three *interpretable* parameters [23]: the mean rate m, equivalently the resource utilization, measuring the *volume* or "quantity" of traffic, and two parameters related to the *burstiness* or "quality" of traffic (the peakedness measure a giving the magnitude of fluctuations about the mean rate, and the Hurst parameter H indicating the rate of decay of correlations in the traffic). All three features have significant impacts on network performance and engineering [9,12,18,22]. The combination of $\{m,\ a,\ H\}$ completely describes the model, and the FBM traffic model is a reasonable representation of *aggregate* data traffic (i.e., formed by multiplexing a large number of independent data sources), as has been observed and validated in traffic analyses of various packet-network technologies and services as noted above. In comparing FBM models to real traffic, the match is poorer at smaller time scales where physical limitations and/or protocol behaviors (e.g., TCP window congestion control mechanism) may govern traffic generation, but the match at intermediate time scales seems to extend to the longest time scales for which the data allow a comparison. These comparisons show that while real traffic is not exactly self-similar, it appears to be asymptotically self-similar.

Bellcore GR-1248-CORE - Generic Requirements for Operations of ATM NEs (referred to as GR-1248 hereafter) specifies the measurements that must be supported by a Broadband Switching System (BSS) either as part of regularly scheduled measurements, or in special studies. These will be reviewed in Section 3. It is anticipated, however, that either the measurements specified in current GR-1248 alone may not be sufficient or the equipment suppliers may choose not to implement measurements specified in GR-1248. Section 4 suggests alternative measurements that can support engineering of bursty traffic, which is the essential feature of most applications on high-speed networks.

3. A REVIEW OF RELEVANT BELLCORE GR-1248 REQUIREMENTS

Section 8 of GR-1248 - *Network Data Collection (NDC) Requirements* describes measurements gathered over an extended period of time to detect violation of service subscription parameters by Permanent

Virtual Circuit (PVC) customers, and to observe longer term trends in traffic patterns and loads.

There are two types of *operational* data collection methods: *scheduled measurements*, generally gathered on a regular and continuous basis to measure usage and monitor the health of the network, and *special studies measurements*, gathered for a limited time period upon Management System's requests for specific traffic characteristics, e.g., to isolate a traffic problem.

Three types of *scheduled* traffic measurements are specified in GR-1248:

- Traffic load measurements: counts of cells entering and leaving ATM Network Element (NE) interfaces and Virtual Path Link (VPL)/ Virtual Circuit Link (VCL). These include: incoming and outgoing cells per interface, incoming and outgoing cells per VPL, incoming and outgoing cells per VCL, traffic load of (i.e., cells counts processed by) each NE module (e.g., processor), and traffic load of (i.e., incoming and outgoing cell counts per) internal link connecting ATM NE modules. ATM NE modules refer to NE traffic-sensitive components that thus need to be traffic engineered.

- Usage Parameter Control (UPC)/ Network Parameter Control (NPC) disagreement measurements: measures of cells discarded due to UPC/NPC violations, and gathered per VPL/VCL.

- Congestion measurements: counts of cells discarded by ATM NE modules due to buffer overflows or intentional discards made to alleviate congestions, as well as the amount of time a congestable resource is in the different levels of congestion. The cell discard measures do not include cells discarded due to protocol nonconformance.

Special studies measurements, used to obtain detailed information on selected traffic, can be used to analyze Operation and Maintenance (OAM) traffic separately from aggregate (i.e., user information + OAM) traffic, and to examine aggregate traffic over shorter time intervals to study traffic behavior. For NDC, two types of *special studies* measurements are specified:

- 15-minute special studies counts: snapshots of 15-minute time intervals for OAM cell count analysis, basically performed to determine whether the volume of OAM traffic is limited to the small number of cells that it should be. This capability should be built into the switch.

- 1-second special studies counts: snapshots of 1-second time intervals for traffic characterization. These can be started at the beginning of the next 15-minute interval, and extend for multiples of 900 seconds (15 minutes). This capability may be built into the switch or into external equipment.

In terms of the statistical features relevant to ATM performance discussed in Section 2, the scheduled ATM traffic measurements provide average rates over 15 minute intervals, too coarse to capture the variations within finer time scales (e.g., time scales relevant in queuing [22]); these alone are not adequate for traffic engineering. The next section discusses how relevant traffic parameters and engineering/performance measures can be estimated from other traffic measurement alternatives.

4. MEASUREMENT ALTERNATIVES; THEIR USAGE IN TRAFFIC ENGINEERING

Given implementation and capacity costs of traffic measurements, proposed measurements must add significant value to the traffic engineering process. We thus consider how various measurements can be useful in practice. The objective is to capture as much information as possible at a small fraction of overhead. Where possible, demonstrations of the usage and usefulness of measurements are performed based on high-resolution traces collected from a variety of high-speed networks including Ethernet [19], Frame Relay [15], ATM [16,17], and synthetically generated FBM streams [24]. These examples are illustrative; more comprehensive experiments are required to make conclusive inferences regarding the use of these measurement alternatives.

4.1. Special Studies Cell Counts

In this measurement, the BSS or adjunct equipment collects cell counts over a small time interval, and reports the entire time series of cell counts at the end of the engineering period. An example of this is specified in GR-1248, which states that the BSS, possibly in conjunction with test/adjunct equipment, should support one second cell count special study within a data collection interval of 15 minutes, resulting in a time series of 900 cell counts. Another example is cell counts per frame in VBR video. In either case, the time scale over which cell counts are accumulated is relatively coarse compared to cell emission times. There is clearly considerable loss of information with this measurement in that all fluctuations in traffic on time scales shorter than this time scale are lost; but there is substantial reduction in overhead compared to that required for in collecting, storing and transporting complete information.

If the traffic is known to be self-similar or fractal (e.g., LAN interconnection traffic), in principle special study cell counts can be used to estimate both the fluctuations in magnitude (the peakedness factor a), as well as the correlation parameter H, [2] since in theory the same values of traffic parameters are obtained whether the traffic measurements are based on 1 millisecond (ms) counts or 1 second (sec) counts. In practice, time series of cell counts on the order of 900 counts (or even 3600 counts in a busy hour) may not be sufficient to estimate traffic parameters *accurately*. But if there is prior information on these parameters (e.g., based on an earlier high resolution traffic study) special studies can validate the choices of these parameters, or indicate if underlying traffic characteristics have changed.

One method which can use cell counts from smaller intervals within a (e.g., 15 minute) recording period is the wavelet-based analysis [1,2]. Figure 2 illustrates this estimation method for each of the Frame Relay, Ethernet, ATM and FBM datasets mentioned in the Introduction, based on 5-sec, 1-sec, and 100-ms cell counts over a busy hour. Table 1 shows the Hurst estimates obtained from these plots plus from a time series of 10-sec counts. Estimations using the 100 ms counts are the ones considered most accurate for engineering purposes, since the time periods of interest for queueing, although dependent on the network system under study, generally fall between 50 ms and 500 ms. As seen, as the measurement interval increases beyond a few seconds, the estimates become unreliable (with a large confidence interval) mainly due to insufficient number of measurements for making meaningful statistical inference.

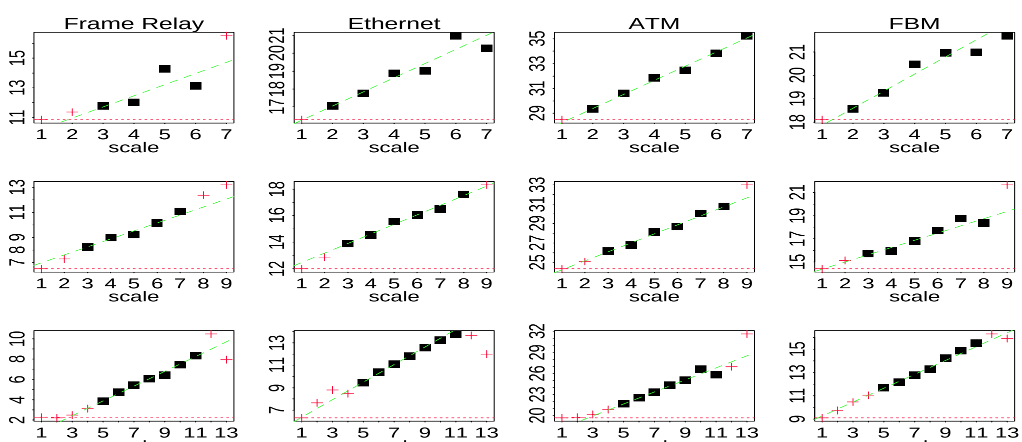

Figure 2. Hurst parameter estimates from cell counts over 3 time scales: 5, 1, and 0.1 seconds; the darker points were used in estimation

Experience with supplier implementations indicates that even a 1-sec data measurement interval may be too frequent and costly, so different measurement strategies may be required. These can include:

1. coaser time scale measurements (e.g., 10 sec) within a reporting period, but with collections performed for the same reporting period each day (e.g., one busy hour per day) in order to obtain

[2] By "self similarity" we mean that the traffic process has the same underlying statistical characteristics on whichever time scale it is measured.

	Frame Relay	Ethernet	ATM	FBM
10 seconds	$1.10 \pm .23$	$0.92 \pm .14$	$1.05 \pm .14$	$0.87 \pm .14$
5 seconds	$0.88 \pm .17$	$0.90 \pm .09$	$1.07 \pm .09$	$0.87 \pm .09$
1 second	$0.82 \pm .06$	$0.86 \pm .06$	$0.97 \pm .06$	$0.81 \pm .06$
100 ms	$0.86 \pm .03$	$0.89 \pm .03$	$0.93 \pm .03$	$0.80 \pm .03$

Table 1
Hurst parameter estimates from cell counts over different time scales

a larger number of samples and therefore increase confidence in the estimate. In this case, an adjustment factor for the variances can be derived similar to that in the Appendix. Assume that n (correlated) measurements are taken in a busy hour for d independent days, then the variance adjustment factor in this case can be shown to be (derivations omitted here due to page limitation)

$$\gamma' = dn - n^{2H-1} \tag{1}$$

Note that when $d = 1$, $\gamma' = n(1 - n^{-2(1-H)}) = \gamma$ in Equation (5). Using this adjustment factor, the Hurst parameter can be estimated through the iterative approach described earlier.

2. nonhomogeneous sampling, where the engineering period is divided into non-overlapping intervals of different time scales (10 ms, 100 ms, 1 sec, etc.) from which measurements are taken. For example, it may be advantageous to take more samples at both ends of the engineering period and less frequent samples in the middle due to the traffic long-range dependence. Further research is required in this area.

4.2. Variance-Time Measurements

Another measurement currently supported by some lower speed X.25 packet switching vendors is sums of squares of frame or byte counts over fine time scales. While this is an indication of the magnitude of underlying traffic fluctuations, this type of measurement will not in itself be sufficient in ATM traffic engineering. But if such measurements are available over a *range* of time scales, in theory the correlation structure underlying the traffic stream and the parameters impacting performance can be estimated. Such measurements are suggested as an alternative to reporting finer scale time series of cell counts since the number of reporting measurements is significantly less. These measurements are not currently specified in GR-1248, but their use is suggested if adjunct equipment or the BSS itself is capable of reporting them. Since some low speed packet switch vendors do in fact provide such measurements, they are technically and possibly economically implementable.

Given the sums of squares of cell counts over a range of intervals (e.g., 10 ms, 50 ms, 100 ms, 1 sec, 5 sec, 10 sec, 50 sec), and the regularly scheduled total interval average cell counts (e.g., 15 minutes), the variance of counts over these time scales can be readily estimated. A log-log plot of variance vs. time scale can be used to estimate the Hurst parameter H, as well as other correlation coefficients [19]. Figure 3 depicts this method; the derived parameter estimates are provided in Table 2. Except for ATM which underestimates H, the H estimates are in line with those provided by the wavelet procedure.

Note that for some scenarios, the variance-time measurements may not be as clean as the measurements shown in Figure 3. As mentioned at the beginning of this section, the examples provided here are not sufficient for inference regarding the use of the alternative measurement.

	Frame Relay	Ethernet	ATM	FBM
H (10 ms to 50 sec variances)	0.83	0.85	0.84	0.83

Table 2
Hurst parameter estimates from variances obtained from sums-of-squares measurements

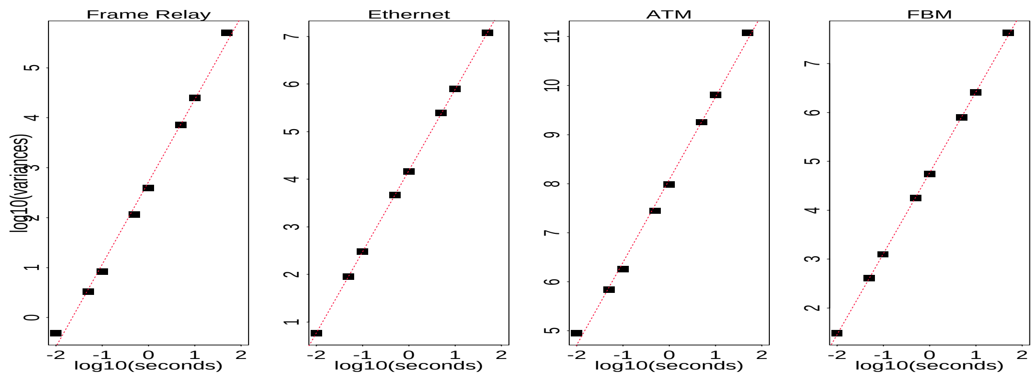

Figure 3. Estimation of H using variances obtained from sums-of-squares measurements

4.3. Peak Rate Measurements

In this class of measurements, traffic counts or utilization measurements are made on a fine time scale, but collection overhead is reduced by reporting only the nth highest value. An arbitrary n increases measurement overhead; we restrict ourselves here to the case when $n = 1$. The practical utility of reporting the n-th peak value will be considered in the future. Peak rates are not currently included in GR-1248, but reporting the peak value, or high water mark, is feasible in practice, judging by the fact that such measurements are available in a number of low speed packet switches.

4.3.1. Using PMR for Predicting Capacity Exhaust

Busiest intervals can dominate even average measures of performance. Simulations show that peak rates over a carefully chosen interval can convey valuable information on capacity exhaust [10]. Figure 4 shows a load-service curve obtained through simulation using an actual trace. As load levels increase beyond the knee of the curve, seen around 30% on the left panel of Figure 4, a small increase in load can result in substantial degradation in performance. For operating points below the knee, utilizations and network efficiencies can be improved without incurring a performance penalty.

The peak-to-mean ratio (PMR) is obtained by dividing the peak cell count by the average count over the engineering period. If the average utilization over the engineering period is U%, a PMR of k indicates that the utilization is $k \cdot U$ on finer time scales. A $k \cdot U$ consistently approaching the maximum possible utilization (100%) indicates that the resource is getting congested over the smaller time scale. That is, the capacity exhaust should occur when the long term utilization level is at about 1/PMR. For example, if PMR $= 1.751$, then the capacity exhaust should occur at utilization of 0.571. The length of the interval over which the peak rate is measured is crucial. For intervals too short, the high water mark can be extremely peaked, and setting operating points on this basis may be overly conservative. For intervals too long, traffic peaks can get smoothed over, with the result that the operating point can be unduly optimistic. This is illustrated by comparing the capacity exhaust predicted by a range of 1/PMRs with an actual load-service curve. Figure 4 depicts average waiting times (in ms) against the utilization level for Frame Relay (left), Ethernet (middle), and ATM (right) based on single server simulation analyses with actual Ethernet, Frame Relay, and ATM traces as the input. The five vertical lines, from left to right, denote 1/PMR measured over 10 ms, 100 ms, 1 sec, 10 sec and 100 sec.

The relevant time scale of queuing interest has been shown to be [22]

$$t_q = \frac{B}{C-m} \cdot \frac{H}{1-H} \tag{2}$$

where B is buffer size, C is resource (e.g., link) capacity, m is average arrival rate, and H is the Hurst parameter. For Ethernet LAN interconnection services (with assumptions of buffer sizes in the range of 100 to 1000 frames, mean rate of 10%, and H between 0.7 and 0.9), the peak rates measured over intervals of the time scales (t_q) of 100 ms to 1 sec can reasonably indicate exhaust (also indicated by our earlier experimental studies [10]), while for Frame Relay, the measurement interval should roughly be in the range of 1 sec to 10 sec. Figure 4 confirms that the peak-to-mean ratio measured over the relevant queuing time scale can be used to indicate the knee of the load-service curve and thus the threshold of capacity exhaust. As this is dependent on various factors (see Equation (2)), the network provider should have the flexibility to set the measurement interval over a range of time scales.

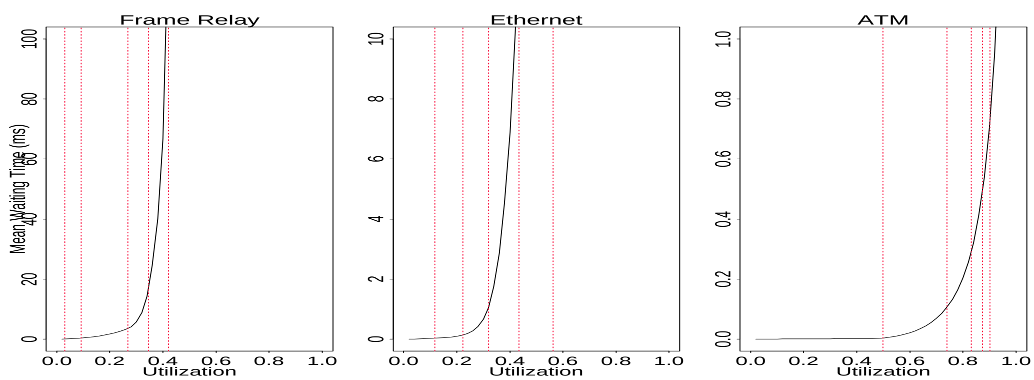

Figure 4. Using PMR ratio to predict capacity exhaust

While PMRs over a carefully chosen interval can indicate capacity exhaust, they can not (in particular, in scenarios where σ is not much greater than the mean rate) be used as inputs to models that address other engineering concerns, such as predicting multiplexing gains. Further, engineering purely on the basis of peak levels may be very uneconomical. The use of the PMR is only suggested here as an interim measure, when more detailed measurements indicating the correlation behavior are not feasible.

4.3.2. Using Peak Measurement for Traffic Parameter Estimation

In *certain scenarios*, the peak rate measurement can be used to estimate the Hurst parameter based on the following scaling relation (derivations omitted here). Let $\hat{P}(t)$ denote the peak measurement at time scale t,

$$\hat{P}(t_1) = \hat{P}(t_2)\left(\frac{t_1}{t_2}\right)^H \tag{3}$$

Figure 5 illustrates the method, and the derived H parameter estimates (for the four datasets mentioned previously) are provided in Table 3. The estimates for the Frame Relay and the FBM datasets are not in close agreement with those from the wavelet estimations, while those for Ethernet and ATM compare well with those from the wavelet estimations. This illustrates the difficulty in using peak rates in traffic engineering, which lies in the interpretation of a statistic that is not well behaved.

4.4. Link and Processor Utilization Histograms

Resource average utilization over an engineering period may not be sufficient to represent actual usage in bursty traffic environments. Some low speed packet switching vendors provide more information in terms of a utilization histogram, a report of the percentage of time (for finer time scales) usage of the network resource was within a range of utilization values. Measures of congestion required by GR-1248

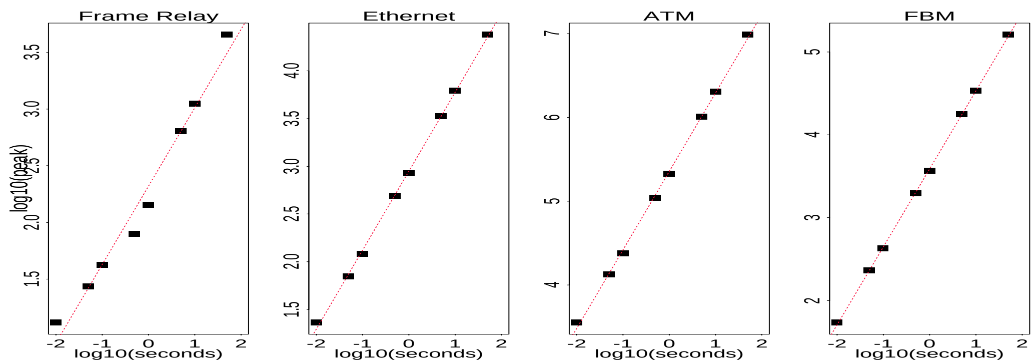

Figure 5. Estimation of H using scaling relationship with peak measurements

	Frame Relay	Ethernet	ATM	FBM
H (10 ms to 50 sec peak)	*0.69	0.83	0.94	*0.94

*Estimates which appear out of alignment with estimates from other methods.

Table 3
Hurst parameter estimates from peak counts

are another example of this measurement, with four utilization values. Figure 6 provides examples of utilization histograms with ten bins, with bin size 10%; indicated is the percentage of time the resource spends within a given range. The Ethernet network resource had an average utilization of 18%, but 5% of the time was utilized more than 50%. Once again this measurement is based on observing the traffic over smaller time scales (e.g., GR-1248 specifies a 20 ms time scale), but the collection overhead is reduced by reporting a set of summary statistics (the histogram) at each reporting period.

In principle, a utilization histogram provides more information than the PMR, but as with PMRs its practical utility is determined by the underlying fine time scale measurement interval. Since performance is dominated by the busiest intervals, much histogram information is not particularly useful beyond a qualitative understanding of traffic burstiness. While such measurements are important in congestion control, their use in engineering is less clear. It may not be possible to estimate the traffic parameters described in Section 2 on the basis of link or processor utilization histograms; e.g., correlation depends on the time sequence of the usage measurements, not just their magnitudes as seen in a histogram.

4.5. Buffer Utilization Measurements

The scope of the measures of congestion measurement in GR-1284 also includes buffer usage measurements which are more useful than link or processor utilization histograms. First, buffer usage measurements (either averages, or histograms indicating the number of time intervals spent in various usage levels) are a direct indication of the impact of the traffic on the network element, and consequently, of performance seen by customers. In conjunction with measurements describing traffic, buffer usage monitoring can be used: 1) to validate engineering rules; 2) in Connection Admission Control (CAC) algorithms [11]; and 3) in congestion controls.

Buffer usage measurements, interpreted as backlogs in a queueing system, can also be used to predict capacity exhaust by the following two-step procedure:

1. During initial deployment stages, characteristics of underlying traffic patterns are estimated from measurements of rates and buffer usage.

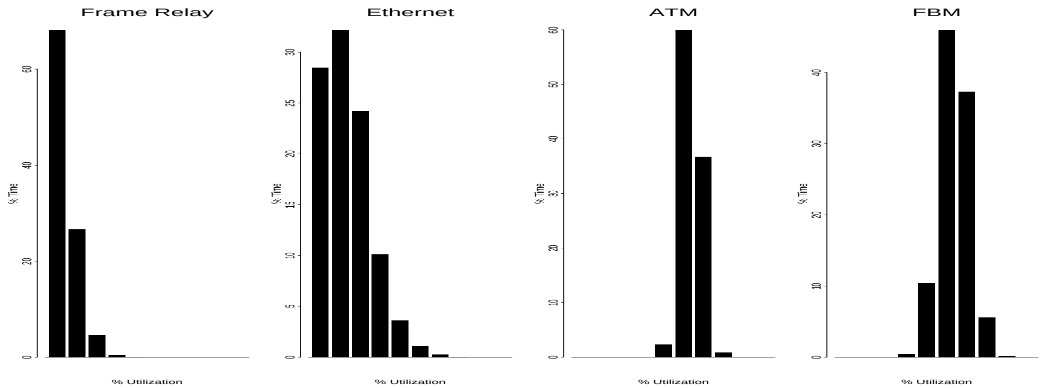

Figure 6. Examples of utilization histograms

2. Based on estimated characteristics, a model can extrapolate effects of increasing traffic levels, and estimate the utilization level at which performance degrades and capacity exhausts.

On-going buffer usage measurements can be used to continually refine the capacity exhaust model. Given a suitable model, it is expected that such an approach can yield more accurate indications of capacity exhaust and safe network operating points. Therefore, the BSS should report, as part of regularly scheduled measurements or as part of special studies, some measure of the usage of buffer (input, output or switch fabric) resources: average buffer usage; percentage of time usage exceeds a threshold; or histogram of actual usage, i.e., the percentage of time buffer usage is in a given range. Average buffer usage can be estimated on the basis of periodic scans that poll the instantaneous buffer backlog. To reduce the variability of the measurement, the scans should be done over fine time scales, though the average buffer usage measurement may be reported only every 15 minutes.

5. SUMMARY

We reviewed and proposed several operational traffic measurements, those specified in Bellcore generic requirements (GR-1248) and new measurement alternatives, and described how these measurements can be used in broadband ATM network engineering. The new measurement alternatives include:

- Time series of cell counts over finer time scales can be collected as part of special studies, and used, under some conditions, to estimate relevant traffic parameters.

- Variance-time measurements are a good alternative to reporting time series of cell counts. If adjunct equipment with the capability to collect sums of squares of cell counts is available, these measures should be obtained over a range of time scales; useful cell count intervals can include 100 microseconds, 1 ms, 10 ms, 100 ms, 1 sec, 10 sec and 100 sec reported over 15-minute intervals. In principle, such measurements can indicate the correlation structure underlying traffic streams, and provide the information necessary to address most traffic engineering issues.

- Peak rate measurements can be an interim measure if more detailed measurements for traffic correlation behavior are not available. Peak-to-mean ratios over carefully selected time intervals (i.e., the queuing time scale) can reasonably indicate capacity exhaust, but since peak level measurements are not well-behaved statistics, they are not adequate to address other engineering issues.

- Link and processor utilization histograms provide a qualitative indication of the burstiness of the traffic, but do not yield information on all the relevant traffic parameters.

- Buffer usage measurements are an indication of the performance seen by customers. In conjunction with measurements describing traffic, buffer usage monitoring can be used to validate engineering rules, in Connection Admission Control (CAC) algorithms, and in congestion controls.

One or more of these new measurement alternatives can potentially be implemented by the switch or management system suppliers to support high-speed network traffic engineering and capacity expansion on a regular and on-going basis.

REFERENCES

1. P. Abry and D. Veitch, "Wavelet Analysis of Long-Range Dependent Traffic," *IEEE Trans. Info. Theory*, Vol. 44, No. 1, pp. 2-15, January 1998.
2. P. Abry and D. Veitch, "Long-Range Dependence: Revisiting Aggregation with Wavelets," *Journal of Time Series Analysis*, Vol. 19, No. 3, pp. 253-266, May 1998.
3. "cflowd," `http://www.caida.org/Tools/Cflowd`.
4. M. Roughan, D. Veitch and P. Abry, "On-Line Estimation of Long-Range Dependence Parameters," *Proc. IEEE Globecom*, pp. 3716-3721, Sydney, Australia, November 1998.
5. J. Beran, R. Sherman, M.S. Taqqu and W. Willinger, "Long-Range Dependence in Variable-Bit-Rate Video Traffic," *IEEE Trans. Commun.*, pp. 1566-1579, 1995.
6. M.E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *Proc. ACM Sigmetrics*, pp. 160-169, May 1996.
7. D.E. Duffy, A.A. McIntosh, M. Rosenstein and W. Willinger, "Statistical Analysis of CCSN/SS7 Traffic Data from Working Subnetworks," *IEEE JSAC*, pp. 544-551, 1994.
8. A. Erramilli, L.J. Forys and J.L. Wang, "New Traffic Engineering and Analysis Methods for Emerging Technologies," *Proc. IEEE Globecom*, pp. 848-854, Singapore, November 1995.
9. A. Erramilli, O. Narayan, W. Willinger, "Experimental Queueing Analysis with Long-Range Dependent Packet Traffic," *IEEE/ACM Trans. on Networking*, Vol. 4, No. 2, pp. 209-223, April 1996.
10. A. Erramilli and J.L. Wang, "Monitoring Packet Traffic Levels," *Proc. IEEE Globecom*, pp. 274-280, San Francisco, CA, 1994.
11. A. Erramilli and J.L. Wang, "A Connection Admission Control Algorithm for Self-Similar Traffic," submitted for publication, 1999.
12. A. Erramilli, W. Willinger and J.L. Wang, "Modeling and Management of Self-Similar Traffic Flows in High-Speed Networks," to appear *State-of-the-Art in Performance Modeling and Simulation*, K. Bagchi (ed.), Gordon and Breach.
13. M.W. Garrett and W. Willinger, "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic," *Proc. ACM Sigcomm*, pp. 269-280, London, UK, 1994.
14. *Generic Requirements for Operations of ATM NEs*, Bellcore Generic Requirements GR-1248-CORE, Issue 3, August 1996.
15. J.L. Jerkins, J. Monroe, M. Pucci and J.L. Wang, "Carrying PSTN-Offloaded Internet Traffic Over Frame Relay: Frame and Call Level Traffic Analyses," to appear in *Proc. NetWorld+Interop Engineers Conference*, Las Vegas, NV, May 1999.
16. J.L. Jerkins and J.L. Wang, "A Measurement Analysis of ATM Cell-Level Aggregate Traffic," *Proc. IEEE Globecom*, pp. 1589-1595, Phoenix, AZ, November 1997.
17. J.L. Jerkins and J.L. Wang, "Establishing Broadband Application Signatures through ATM Network Traffic Measurement Analyses," *Proc. IEEE ICC*, pp. 837-843, Atlanta, GA, June 1998.
18. K.R. Krishnan, A.L. Neidhardt and A. Erramilli, "Scaling Analysis in Traffic Management of Self-Similar Process," *Proc. 15th ITC*, pp. 1087-1096, Washington, DC, June 1997.
19. W.E. Leland, M.S. Taqqu, W. Willinger and D.V. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Trans. on Networking*, Vol. 2, No. 1, pp. 1-15, 1994.
20. K. Meier-Hellstern, P.E. Wirth, Y.-L. Yan and D.A. Hoeflin, "Traffic Models for ISDN Data Users: Office Automation Application," *Proc. 13th ITC*, pp. 167-172, Copenhagen, 1991.
21. "Netflow Services and Applications," Cisco IOS Software Solutions, `http://www.cisco.com/warp/public/732/netflow`.
22. A.L. Neidhardt and J.L. Wang, "The Concept of Relevant Time Scales and Its Application to Queuing Analysis of Self-Similar Traffic (or Is Hurst Naughty or Nice?)," *Proc. ACM Sigmetrics*, pp. 222-232, Madison, WI, June 1998.
23. I. Norros, "A Storage Model with Self-Similar Input," *Queueing Systems*, Vol. 16, pp. 387-396, 1994.

24. I. Norros, P. Mannersalo and J.L. Wang, "Simulation of Fractional Brownian Motion with Conditionalized Random Midpoint Displacement," to appear in *Advances in Performance Analysis*, 1999.
25. V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," *Proc. ACM Sigcomm*, pp. 257-268, London, UK, 1994.

APPENDIX: REVISED VARIANCE-TIME ANALYSIS

The variance-time analysis is a popular method for detection of long-range dependence (LRD) and estimation of the LRD parameter, i.e., Hurst parameter. Consider $X = (X_t : t = 1, 2, \cdots)$ a covariance stationary stochastic process with mean $\bar{X}$ and variance σ^2 and autocorrelation function $r(k),\ k \geq 0$. Let $X^{(m)} = ({X_k}^{(m)} : k = 1, 2, \cdots),\ m = 1, 2, \cdots$, denote new covariance stationary processes (with autocorrelation function $r^{(m)}(k)$) obtained by averaging the original series X over non-overlapping blocks of size m; that is, $X^{(m)}$ is given by ${X_k}^{(m)} = 1/m(X_{km-m+1} + \cdots + X_{km}),\ k \geq 1$. Then, one evidence of X being long-range dependent is that

$$V(X^{(m)}) \sim cm^{-\beta}, \text{as } m \to \infty \tag{4}$$

with $0 < \beta < 1$, where $V(X)$ denotes the variance of X and c is a positive constant. β is related to the Hurst parameter through $\beta = 2(1 - H)$. If $\beta = 1$, then the time series X is short-range dependent. The so-called *variance-time plots* are obtained by plotting $\log(V(X^{(m)}))$ against $log(m)$ and are used to determine whether or not the time series is LRD and to estimate the Hurst parameter value.

If the original time series X is long-range dependent then the variances of $X^{(m)}$ at various m are correlated, and these correlations should be taken into account in evaluating the variance for a given m (or time scale). Let $A((j-1)t, jt),\ j = 1, 2, \cdots$ denote the number of arrivals measured during the time interval $(j-1)t$ to jt. Let $V(t) := \mathrm{var}(A(0,t))$ and note that $\mathrm{Cov}[A(p,q), A(s,t)] = \frac{1}{2}\{V(t-p) + V(s-q) - V(t-q) - V(s-p)\}$. Let T be the total measurement interval, t is the interval when individual measurements are taken, and $n = T/t$. Introducing $\| X \|_0$ as the standard deviation of X, we have

$$\begin{aligned}
&\mathbf{E}\left[\sum_{j=1}^{T/t} [A((j-1)t, jt) - A(0,T)/(T/t)]^2\right] \\
= &\sum_{j=1}^{T/t} \| A((j-1)t, jt) - A(0,T)/(T/t)\|_0{}^2 \\
= &\sum_{j=1}^{T/t} \left[\| A((j-1)t, jt)^2\|_0{}^2 - 2\frac{t}{T}\mathrm{cov}(A((j-1)t, jt), A(0,T)) + \left(\frac{t}{T}\right)^2 \| A(0,T)^2\|_0{}^2\right] \\
= &\ nV(t) - n^{-1}\sum_{j=1}^{n}[V(T-(j-1)t) + V(jt) - V(T-jt) - V((j-1)t)] + n^{-1}V(T) \\
= &\ nV(t) + n^{-1}V(T) - n^{-1}[V(T) - V(0) + V(T) - V(0)] \\
= &\ nV(t) + n^{-1}V(T) - 2n^{-1}V(T) \\
= &\ nV(t) - n^{-1}n^{2H}V(t) \\
= &\ V(t)\cdot\left[n(1 - n^{-2(1-H)})\right]
\end{aligned} \tag{5}$$

Then $V(t)$ can be adjusted by dividing the expectation on the left hand side of Equation (5) by a factor $\gamma = n(1 - n^{-2(1-H)})$. Note that when $H = 0.5$, $\gamma = n - 1$ which is the usual independent case to get an unbiased estimate of the variance. For $0.5 < H < 1$, the adjustment factor (or the variance) at any given time scale depends on the Hurst parameter H. Therefore, the revised variance-time analysis needs to start with an estimate of H (perhaps without the aforementioned adjustment), then iterate until the target Hurst parameter is the same as the one (input) in the adjustment factor.

This method can be used in the measurement analyses based on special studies cell counts as well as variance-time measurements.